

Performance of Gigabit FDDI

David Game

Kurt Maly

Department of Computer Science

Old Dominion University

Norfolk, Virginia 23529-0162

April 23, 1990

Abstract

Great interest exists in developing high speed protocols which will be able to support data rates at gigabit speeds. Hardware currently exists which can experimentally transmit at data rates exceeding a gigabit per second, but it is not clear as to what types of protocols will provide the best performance.

One possibility is to examine current protocols and their extensibility to these speeds. This paper investigates scaling of FDDI to gigabit speeds. More specifically, delay statistics are included to provide insight as to which parameters (network length, packet length or number of nodes) have the greatest effect on performance.¹

¹This work was supported by CIT grant RF-89-002-01, NASA grant NAG-1-908 and Sun Microsystems grant RF 596043.

1 Introduction

Lasers, optical fiber, and related optics technology have recently redefined the bottleneck in data communications[2,4,6,7] such that the communications channel is no longer the limit to information processing. Instead, the issue has become one of whether or not a computer can generate and/or process information at the rates which are now available. Applications for interactive video such as a surgeon examining a 3-D display of an organ can readily use such high bandwidth, but the computer itself is not able to generate images at a rate which will consume such a high bandwidth for an extended period of time.

Nonetheless, great interest exists in extending the communications capacity. Although a single computer may not be able to use such high rates, a large number of nodes can, and a national research initiative is ongoing in an attempt to develop a gigabit channel for applications such as a national research network[3].

2 FDDI

FDDI[9] is a 100 Mbps fiber optics ring which is commercially available and currently being used primarily as a backbone for internetwork communication. The cost (about \$10,000 per node) is a major factor prohibiting its use in workstations, but this is expected to drop significantly as the product matures. Given its likely widespread use, we investigate in this paper the effect of using a gigabit transmitter in this type of network.

FDDI is fundamentally a token ring network. The distinctive characteristics of the network are its use of fiber optics and associated high data rates, a dual counter rotating ring topology and a token holding timer algorithm to determine the length of time for which a node may hold the token and transmit data. Although FDDI is a dual ring, the second ring is primarily intended to allow for healing in the event of a damaged link[8]. For this reason, only one ring is modelled.

The token holding timer algorithm is one whereby each node keeps a local timer as a means of determining how long it can hold the token for transmission. It is intended to place a bound on access delay for synchronous traffic. Each time the token arrives, the clock is reset. If a sufficiently small amount of time has expired (less than an amount negotiated among the nodes called the target token rotation time, TTRT), data may be transmitted for TTRT minus the lapsed time on the timer when the token returned. At that point the token is released[1]. Consider the case where the TTRT value is set precisely at the level which will let every node transmit its data on each cycle(rotation) of the token. If TTRT is reduced by one-half, half of the nodes (actually less) will be able to transmit during each rotation. The overhead of passing the token becomes more significant and utilization is decreased. For a more detailed discussion, see [1]. In order to minimize this as

a factor, the TTRT value was set arbitrarily high (20 milliseconds) in these runs. Only asynchronous traffic is considered.

3 Parameters and Metrics

Clearly, extending the rates will improve performance over standard FDDI. Packet transmission times will be proportionally reduced and propagation delays will remain the same. The question is to determine which factors will have the greatest impact on such a network so that the environment in which it can best be utilized can be better understood.

It is anticipated that the predominant factors which will affect performance are

1. number of nodes,
2. length of the network, and
3. packet length.

For the simulation, each of the three parameters above have been tested over a range of three values each as follows.

Nodes	10, 100, 1000
Network Length	1Km, 10Km, 100Km
Packet Length	5K, 10K, 15K

Given the large bandwidth of the network, it is anticipated that large numbers of nodes can be supported. Length has been considered to include LAN, MAN and WAN scenarios and packet length varies from 5000 to 15000 bits.

The metric used to evaluate performance is delay. This is a measure of the time between arrival of the message at the node to delivery of the last bit of the message at the destination. Other metrics frequently used in network analysis include access delay (time to beginning of message transmission), throughput and fairness. Access delay is not graphed because we are concerned with the impact of distance on the network and want to include the impact of propagation delays as the distance is lengthened. In the results shown here, the only cases considered are those in which the network is less than fully loaded so that throughput is equal to offered load. Fairness of FDDI has been shown in [5]. There is no reason to assume that an increased transmission rate will affect fairness so it too has been ignored here.

4 Results

For the purpose of comparison, each of the following graphs use the same scaling. The x-axis represents load on the system in percent of the transmission rate of the

network. The y-axis represents delay in thousands of microseconds. The results of selected runs from the set of runs described previously are shown.

4.1 Nodes

In a typical token ring network, the number of nodes affects performance in two primary areas. First, there exists a delay introduced on the ring at each node which is using the network. Second, for each node capturing the token on a cycle around the ring, the token arrival to other nodes is delayed by an additional token retransmission. In addition, there is a delay between recognition of the token and transmission of the queued packet. This is explained in [1]. For these reasons the number of nodes has an adverse effect on performance, but the impact of number of nodes varies as explained below.

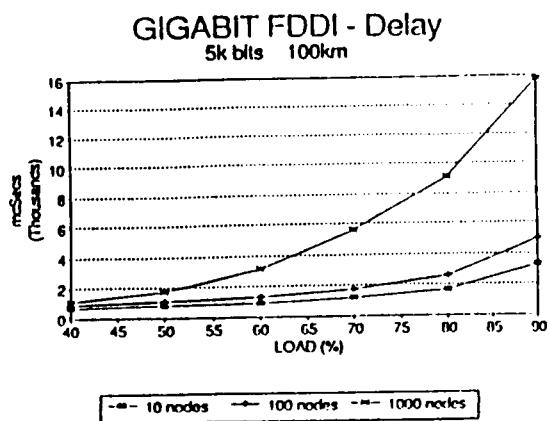
Figure 1 shows the effect of varying the number of nodes in four different scenarios. Vertically the graphs have the same packets size and horizontally they have the same distance. Note that in every case, the number of nodes has a negative effect, however the effect varies with certain combinations of the other parameters. A comparison of the graphs horizontally shows that if the load is distributed in smaller packets, the number of nodes has a greater effect than in if the load the packet size is larger. This can be explained by the fact that the overhead time required for a token capture does have an impact. As the packet size is smaller and thus distributed to more nodes, additional nodes capture the token on each cycle, introducing additional delays.

4.2 Packet Length

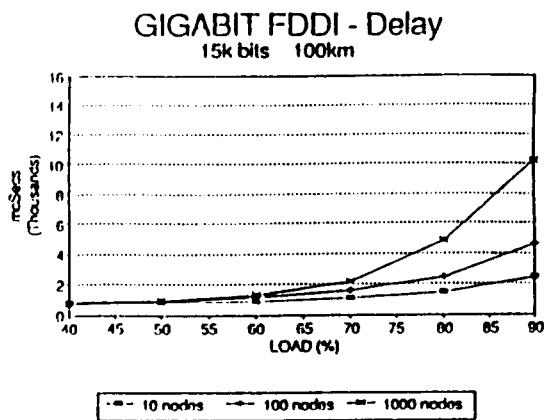
As described in the previous section, packet length and number of nodes, in combination, can have an effect on performance. Figure 2 reinforces the previous results. The three graphs show scenarios where the number of nodes equals 10, 100 and 1000. Notice that in cases a and b, the effect of packet size is practically insignificant. However, when the number of nodes increases to 1000, the delay varies significantly. Case c shows that increasing the packet size from 5000 bits to 15000 bits cuts the delay in half for up to 80% and by a significant quantity for 90%.

4.3 Network Length

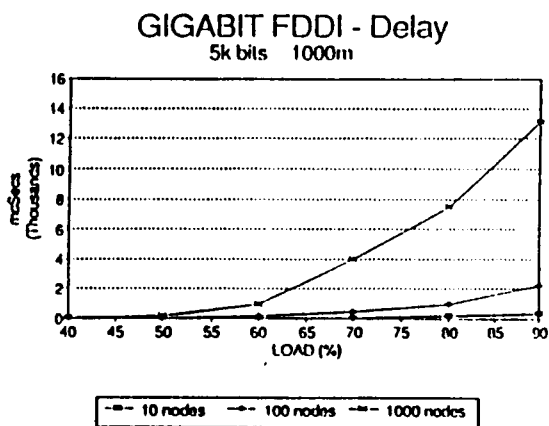
The last set of graphs in Figure 3 shows four scenarios similar to figure 1. One would likely anticipate that the impact of propagation delay is simply a matter of being relatively large for the 100Km case and proportionally less for the other two cases. Cases a and b indicate that although the increased length has a negative effect on the network, it has a worse effect as the number of nodes increase in conjunction with



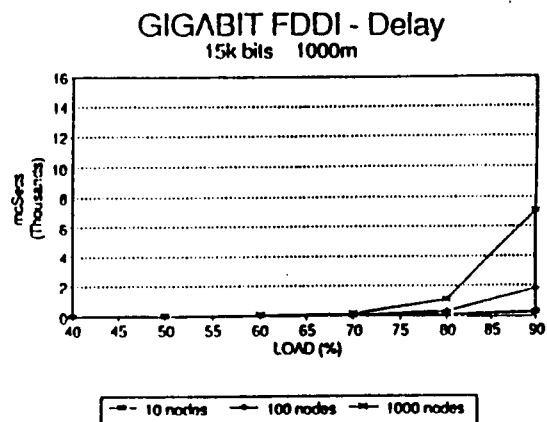
(A)



(B)

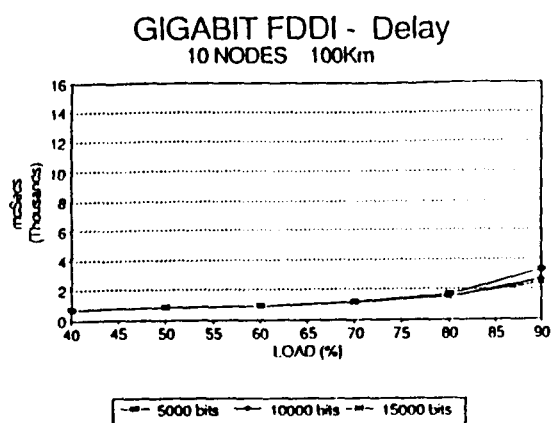


(C)

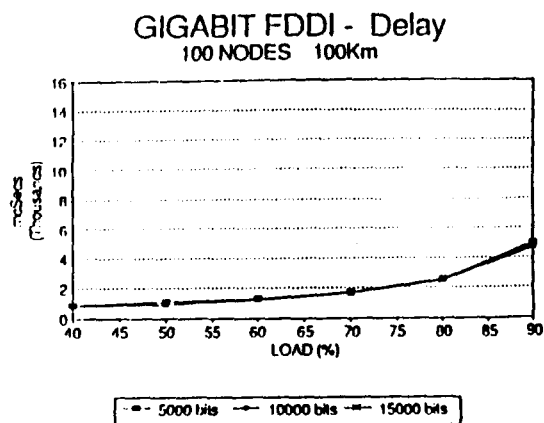


(D)

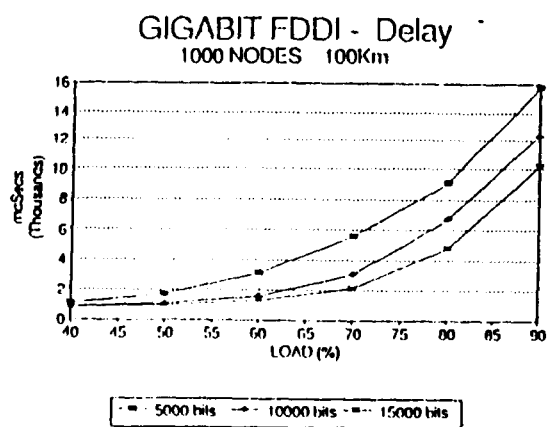
Figure 1: Impact of Nodes



(A)

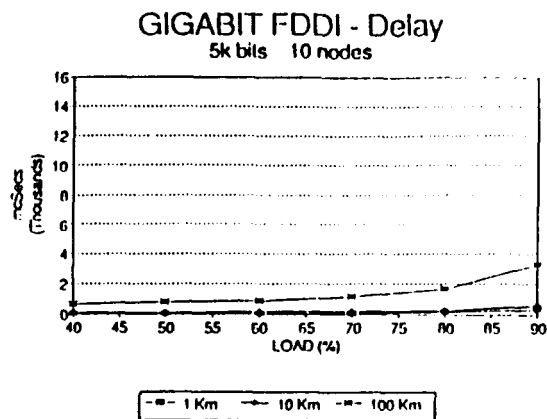


(B)

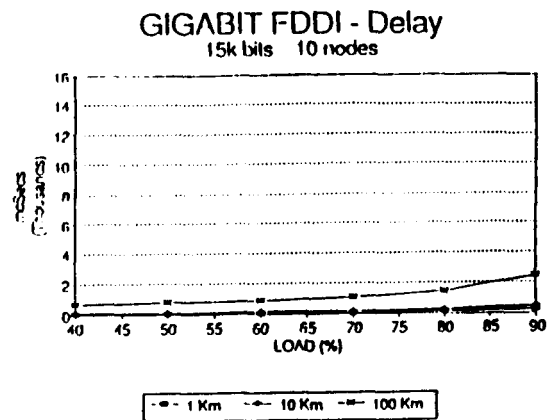


(C)

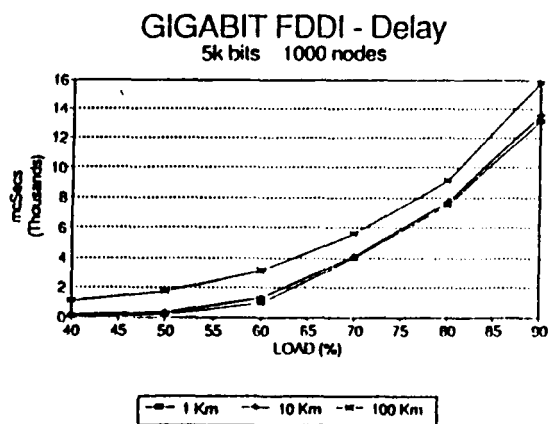
Figure 2: Impact of Packet Size



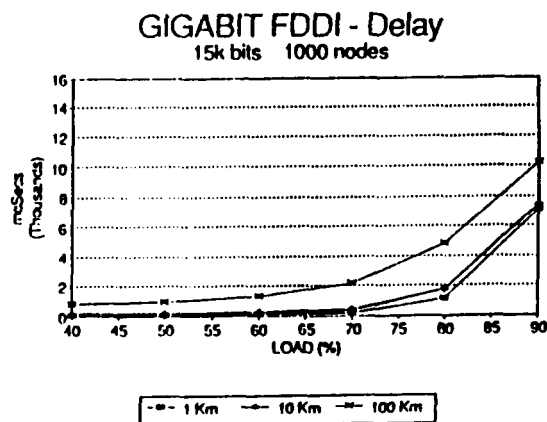
(A)



(B)



(C)



(D)

Figure 3: Impact of Network Length

the network length than if the packet size is decreased in conjunction with the increased network length. In the cases a and b, delay is at or below 2 milliseconds for all values of distance over the entire range of loads.

5 Conclusions

This paper shows the effect of the number of nodes, network length and packet length for FDDI at gigabit speeds. Most of the results show that over the range of parameters examined, delay is on the order of a couple of milliseconds for loads below the 60% level. As load increases above 60%, the delay degrades at different rates depending on the specific case examined. The number of nodes is the one factor which has the greatest effect on performance of the three parameters considered. In addition, the number of nodes compounds the problems worse when increased in conjunction with reducing packet sizes.

Recent research has shown that the number of nodes can in fact be used to reduce the delay and increase throughput through a modification to FDDI. The reader is referred to [1]. Further research should investigate the degree to which the advantages of increased numbers of nodes in a modified FDDI can balance the disadvantages mentioned above and how to what extent increasing packet size will have an advantageous effect on performance.

References

- [1] D. Game and K. Maly. *Extensibility and Limitations of FDDI*. Technical Report, Old Dominion University, February 1990.
- [2] H.S. Hinton. Architectural considerations for photonic switching networks. *IEEE Transactions on Communications*, 6(7):1209-1226, August 1988.
- [3] R.E. Kahn. A national network: today's reality, tomorrow's vision, part 2. *EDUCOM Bulletin*, 14:21, Summer/Fall 1988.
- [4] M. Maeda and H. Nakano. Integrated optoelectronics for optical transmission systems. *IEEE Communications*, 26(5):45-51, May 1988.
- [5] K. Maly, D. Game, L. Zhang, E. Foudrial, and S. Khanna. *Fairness Problems at the MAC Level for High-Speed Networks*. Technical Report, Old Dominion University, February 1990.
- [6] T. Nakagami and T. Sakurai. Optical and optoelectronic devices for optical fiber transmission systems. *IEEE Communications*, 26(1):28-33, January 1988.
- [7] K. Nosu. Advanced coherent lightwave technologies. *IEEE Communications*, 26(2):15-21, February 1988.
- [8] F. Ross. Fddi - a tutorial. *IEEE Communications*, 24(5):10-17, May 1986.
- [9] Draft Proposed American National Standard. Fddi token ring media access control (mac) asc x3t9.5 rev. 10. February 28, 1986.